

Docket Number: POU919980103US1

Inventor: T. A. Gregg

Title: Communicating Between Zones Of A
Central Processing Complex

APPLICATION FOR UNITED STATES
LETTERS PATENT

"Express Mail" Mailing Label No.: EK830786432US
Date of Deposit: December 19, 2000

I hereby certify that this paper is being deposited with the United States Postal Service as "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, DC 20231.

Name: Sandra L. Kilmer

Signature: Sandra L. Kilmer

INTERNATIONAL BUSINESS MACHINES CORPORATION

COMMUNICATING BETWEEN ZONES OF A CENTRAL PROCESSING COMPLEX

Technical Field

The present invention relates, in general, to data
5 communication, and in particular to, moving data between
zones of a central processing complex.

Background of the Invention

A central processing complex (CPC), offered by
International Business Machines Corporation, is capable of
10 concurrently executing a plurality of control zones. For
example, a central processing complex 100 (FIG. 1a) executes
one or more operating system zones 102 and a coupling
facility zone 104. Each operating system zone 102 executes
an operating system, such as for instance, the OS/390 (or
15 Multiple Virtual Storage (MVS)) operating system, offered by
International Business Machines Corporation; and coupling
facility zone 104 executes coupling facility control code
used in aiding the clustering of zones and/or multiple
central processing complexes.

20 Previously, when an operating system zone and a coupling
facility zone within a single central processing complex
wished to communicate with each other, there were two
options available, as shown in FIGs. 1a-1b. In FIG. 1a, an
operating system zone 102 communicated with a coupling
25 facility zone 104 via a pair of I/O channels 106, 108,

interconnected by an external cable 110. In particular, the data is passed from the main memory of the operating system zone to one of the I/O channels of the pair. Thereafter, it is moved from the one I/O channel to its corresponding I/O
5 channel pair, and then to the main memory of the coupling facility zone. Data is similarly moved from the coupling facility zone to an operating system zone. The requirement of a pair of I/O channels, however, disadvantageously adds cost to CPC 100. Further, the I/O channels typically have
10 limited bandwidth.

The second option is described with reference to FIG. 1b. In FIG. 1b, data is moved between an operating system zone 102 and coupling facility zone 104 via one or more processors 120, which execute a coupling facility zone
15 program to move the data. Specifically, special processor instructions are made available to the coupling facility zone that allow its processors to move data between its main memory and main memory of the operating system zone. This internal virtual I/O channel eliminates the need for normal
20 I/O channels, but is costly in that it requires processor resources (i.e., processor instructions) to perform the data movement.

Thus, although various options are currently available for moving data between operating system and coupling
25 facility zones, a need still exists for an improved capability for communicating between zones of a central processing complex. In particular, a need exists for a data movement capability that moves data between the various

zones, and is less costly and has better bandwidth than the previous techniques.

Summary of Invention

5 The shortcomings of the prior art are overcome and additional advantages are provided through the provision of a method of moving data between zones of a central processing complex. The method includes, for instance, initiating a move of data from one zone of the central
10 processing complex to another zone of the central processing complex; and moving the data from the one zone to the another zone without using a channel interface and without using processor instructions.

 In one embodiment, the moving is performed by a data
15 mover located within the central processing complex and coupled to the one zone and the another zone.

 In another embodiment, a method of moving data between zones of a central processing complex is provided. The method includes, for instance, moving data from one zone of
20 the central processing complex to another zone of the central processing complex, wherein the moving includes, for instance: creating one or more queue entries associated with a message requesting the move; generating one or more fetch memory requests for the one or more queue entries to fetch
25 the data from a memory of the one zone; using the one or more fetch memory requests to fetch the data from the memory of the one zone; placing the fetched data in one or more

line buffers; generating one or more store memory requests using the fetched data placed in the one or more line buffers; and employing the one or more store memory requests to store the fetched data in a memory of the another zone.

5 Systems corresponding to the above-summarized methods are also described and claimed herein.

Advantageously, the capabilities of the present invention enable data to be moved between zones of a central processing complex, without requiring a channel interface or
10 processor instructions. This reduces the cost of data movement and provides better bandwidth, without sacrificing protection of memory similarly afforded by channels.

Additional features and advantages are realized through the techniques of the present invention. Other embodiments
15 and aspects of the invention are described in detail herein and are considered a part of the claimed invention.

Brief Description of the Drawings

The subject matter which is regarded as the invention is particularly pointed out and distinctly claimed in the
20 claims at the conclusion of the specification. The foregoing and other objects, features, and advantages of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

FIG. 1a depicts communications between an operating system zone and a coupling facility zone of the same central processing complex using a pair of I/O channels, in accordance with the prior art;

5 FIG. 1b depicts communications between an operating system zone and a coupling facility zone of the same central processing complex using processor instructions, in accordance with the prior art;

10 FIG. 2 depicts one example of a central processing complex including a data mover used in communicating between the zones of a central processing complex, in accordance with an aspect of
15 the present invention;

 FIG. 3 depicts one example of the logical connections between the data mover of FIG. 2 and various zones within a central processing complex, in accordance with an aspect of the present
20 invention;

 FIG. 4 depicts one example of various components of a main memory of an operating system zone and a main memory of a coupling facility zone, in accordance with an aspect of the present
25 invention;

FIGs. 5a-5b depict one embodiment of the logic associated with the flow of data when processing a primary message, in accordance with an aspect of the present invention;

5 FIG. 5c depicts one embodiment of the logic associated with preparing buffers, in accordance with an aspect of the present invention;

10 FIG. 6 depicts one embodiment of the logic associated with the flow of data when processing a secondary message, in accordance with an aspect of the present invention;

FIG. 7 depicts one embodiment of various components of the data mover of FIG. 2, in accordance with an aspect of the present invention;

15 FIG. 8 depicts one embodiment of various components of the queued memory line mover of FIG. 7, in accordance with an aspect of the present invention;

20 FIG. 9a depicts one embodiment of various components of the control logic of FIG. 7, in accordance with an aspect of the present invention;

FIG. 9b depicts one embodiment of a table used by the control logic of FIG. 9a to keep track of

active requests, in accordance with an aspect of the present invention;

FIGs. 10a-10b depict one embodiment of the operation of the control logic associated with processing a primary message, in accordance with an aspect of the present invention; and

FIG. 11 depicts one embodiment of the operation of the control logic associated with processing a secondary message, in accordance with an aspect of the present invention.

Best Mode for Carrying Out the Invention

In accordance with an aspect of the present invention, data is moved directly from one zone of a central processing complex to another zone of the central processing complex without sending the data over a channel interface and without employing processor instructions to perform the move. Instead, a hardware entity (i.e., a data mover) internal to the central processing complex is employed to move the data.

One embodiment of a computing environment incorporating and using aspects of the present invention is described with reference to FIG. 2. Computing environment 200 is based, for instance, on the Enterprise Systems Architecture (ESA)/390 offered by International Business Machines Corporation, Armonk, NY. ESA/390 is described in an IBM POU919980103US1

publication entitled "Enterprise Systems Architecture/390 Principles Of Operation," IBM Publication No. SA22-7201-05, September 1998, which is hereby incorporated herein by reference in its entirety.

5 Computing environment 100 includes at least one central processing complex (CPC) 202, which executes one or more operating system zones 204, and at least one coupling facility zone 206. In this particular example, there are three operating system zones (OS0-OS2) and one coupling
10 facility zone (CF0); however, this is only one example. Each operating system zone executes an operating system, such as, for instance, the OS/390™ (or Multiple Virtual Storage (MVS/ESA™)) operating system offered by International Business Machines Corporation, and each
15 coupling facility zone 206 executes coupling facility control code used in aiding the clustering of zones and/or of multiple central processing complexes.

The coupling facility zone provides a coupling facility, which is a shareable facility that contains storage
20 accessible by the other zones of this complex and other complexes and performs operations requested by a resource manager and/or programs running within the complex(es). In this particular example, the coupling facility is internal to the complex, which is logically partitioned.

25

Aspects of the operation of a coupling facility are described in detail in such references as Elko et al., U.S. Patent No. 5,317,739, entitled "Method and Apparatus For POU919980103US1

Coupling Data Processing Systems", issued May 31, 1994; Elko et al., U.S. Patent No. 5,561,809, entitled "In A Multiprocessing System Having A Coupling Facility, Communicating Messages Between The Processors And The
5 Coupling Facility In Either A Synchronous Operation Or An Asynchronous Operation", issued on October 01, 1996; Elko et al., U.S. Patent No. 5,706,432, entitled "Mechanism For Receiving Messages At A Coupling Facility", issued January 06, 1998; and the patents and applications referred to
10 therein, all of which are hereby incorporated herein by reference in their entirety. Further, aspects of logical partitioning are described in, for instance, Guyette et al., U.S. Patent No. 4,564,903, entitled "Partitioned Multiprocessor Programming System", issued on January 14,
15 1986; Bean et al., U.S. Patent No. 4,843,541, entitled "Logical Resource Partitioning Of A Data Processing System", issued on June 27, 1989; and Kubala, U.S. Patent No. 5,564,040, entitled "Method And Apparatus For Providing A Server Function In A Logically Partitioned Hardware
20 Machine", issued on October 08, 1996, each of which is hereby incorporated herein by reference in its entirety.

Each zone has its own independent area of main memory, and the zones can either have dedicated processor(s) or can logically share processor(s). For example, a single
25 processor can be dedicated to a single zone, or a single processor can be shared by multiple zones.

In accordance with an aspect of the present invention, central processing complex 202 further includes a data mover

POU919980103US1

208 used to move data between the various zones of the CPC.
As one example, data mover 208 moves data between a coupling
facility zone and an operating system zone. As further
examples, data mover 208 may move data between different
5 operating system zones; between different coupling facility
zones; or any combination thereof. The data mover moves the
data between the various zones without requiring a channel
interface between a pair of I/O channels, and without
employing processor resources (e.g., processor instructions)
10 to perform the data movement. The data mover is, for
instance, a hardware entity that automatically moves the
data between the main memories of the zones of the central
processing complex, as described in further detail herein.

In one embodiment, data mover 208 is logically connected
15 to the CPC resources, as shown in FIG. 3. In particular,
data mover 208 has access to each zone's main memory. For
example, data mover 208 has access and is coupled to main
memory 300 of Operating System 0, main memory 302 of
Operating System 1, main memory 304 of Operating System 2,
20 and main memory 306 of Coupling Facility 0. Additionally,
data mover 208 has access and is coupled to a Hardware
System Area (HSA) 308 of main memory. The HSA is a special
area of main memory that is not accessible by any of the
operating system or coupling facility zones. The CPC uses
25 HSA memory for system-wide hypervisor functions, such as
control information.

Each of the connections allows the data mover to move
data directly from one zone's main memory to another zone's

main memory without sending the data over a channel interface or consuming processor resources, as described in detail herein.

In addition to the connections to the main memories, the data mover has a connection to each of the processors of the various zones. For example, data mover 208 is coupled to processors 310, 312, 314, and 316 of Operating System 0, Operating System 1, Operating System 2 and Coupling Facility 0, respectively. Through these connections, each processor can send commands to the data mover and receive status and interruption information.

Further details regarding various components of the main memories of an operating system zone and a coupling facility zone are described with reference to FIG. 4. Additionally, one example of the flow of data when a message is passed from an operating system zone to a coupling facility zone (i.e., a primary message) is described with reference to FIGs. 4 and 5a-5b; and one example of the flow of data when a message is passed from a coupling facility zone to an operating system zone (i.e., a secondary message) is described with reference to FIGs. 4 and 6.

Referring to FIG. 4, an operating system zone's main memory 400 includes, for instance, a Message Control Block (MCB)/Message Response Block (MRB) 402 for storing commands and responses of a primary message; a Link Control Block (LCB) 404 used in sending the command of the primary message to the main memory of a coupling facility zone; one or more

data areas 406, 408 for storing data of the primary message;
an address list 410 for referencing the data of data areas
406, 408; and an LCB 412 used in receiving commands of a
secondary message from the coupling facility zone and
5 sending responses of the secondary message to the coupling
facility zone.

A coupling facility zone's main memory 420 includes, for
instance, a Channel Buffer Operation Block (CBOB) 422 for
receiving commands of primary messages from LCB 404; one or
10 more data areas 424, 426, 426, 430 for receiving data of the
primary message; an address list 432 to reference data areas
424, 426, 428, 430; and a CBOB 434 for moving
responses/commands of a secondary message to/from LCB 412.

One embodiment of the flow of data when a message (i.e.,
15 a primary message) is passed from an operating system zone
to a coupling facility zone is described with reference to
FIGs. 4, and 5a-5b. Referring to FIG. 5a, initially, the
operating system prepares various areas to be used for the
data flow, STEP 500. For example, the operating system
20 builds a message command (MCB) in MCB/MRB area 402 (FIG. 4),
initializes address list 410, and prepares data areas 406,
408 for either a read (data transmitted from the coupling
facility zone to the operating system zone) or a write (data
transmitted from the operating system zone to the coupling
25 facility zone).

Thereafter, the operating system issues a send message
instruction to commence moving the data from the operating

system zone to a coupling facility zone, STEP 502 (FIG. 5a). One example of a send instruction is described in Elko et al., U.S. Patent No. 5,706,432, entitled "Mechanism For Receiving Messages At A Coupling Facility", issued January 5 6, 1998, which is hereby incorporated herein by reference in its entirety.

The processor microcode of the operating system zone issuing the command interprets the send message instruction and moves the MCB from MCB/MRB area 402 (FIG. 4) to Link 10 Control Block (LCB) 404, STEP 504 (FIG. 5a). The microcode then puts the zone ID of the operating system that is executing the send message instruction into the MCB header, STEP 506.

Next, the microcode sends a start hardware command to 15 the data mover instructing it to move the MCB from LCB 404 (FIG. 4) to CBOB 422, STEP 508. One example of a start command is described in detail in Gregg et al., U.S. Patent No. 5,944,797, entitled "Data Mover Hardware Controlled Processing In A Commanding System And In A Commanded System 20 For Controlling Frame Communications On A Link", issued August 31, 1999, which is hereby incorporated herein by reference in its entirety.

When the data mover receives the start command, it checks whether CBOB 422 is ready to receive the command, 25 STEP 509. In one example, it makes this determination by checking a prepared bit in a status buffer of the control logic of the data mover (described below). In particular,

POU919980103US1

in order to prepare the CBOB, the coupling facility zone issues a prepare instruction, STEP 550 (FIG. 5c). (One example of a prepare instruction is described in Elko et al., U.S. Patent No. 5,706,432, entitled "Mechanism For
5 Receiving Messages At A Coupling Facility", issued January 6, 1998, which is hereby incorporated herein by reference in its entirety.) Then, the coupling facility microcode sends a prepare hardware command to the data mover, which includes the address of the CBOB in the coupling facility zone's main
10 memory, STEP 552. The control logic of the data mover sets the CBOB address and sets the prepared state for the message buffer, STEP 554.

If the CBOB is not ready, INQUIRY 510, then the operating system microcode generates and MRB indicating
15 such, STEP 511, and processing of the primary message is complete. However, if the CBOB is ready, then the data mover moves the command to the CBOB, STEP 512.

Additionally, the microcode initiates movement of data areas by the data mover by sending a list hardware command
20 to the data mover when there are data areas to be transferred, STEP 513. One example of a list command is described in Gregg et al., U.S. Patent No. 5,948,060, entitled "Speeding-Up Communication Rates On Links Transferring Data Structures By A Method Of Handling
25 Scatter/Gather Of Storage Blocks In Commanded Computer Systems", issued September 7, 1999, which is hereby incorporated herein by reference in its entirety.

The data mover then waits for the coupling facility zone to prepare its addresses before proceeding with the data transfer, STEP 514. In one example, the data mover detects that the coupling facility is prepared when the data mover
5 has received a list command from both the operating system zone and the coupling facility zone.

In the meantime, when coupling facility zone 420 (FIG. 4) detects the arrival of MCB in CBOB 422, it processes the command, STEP 515. If the command specifies a data
10 transfer, INQUIRY 516, then the coupling facility zone prepares one or more data areas 424, 426, 428, 430, initializes address list 432, and executes an instruction to move the data, STEP 518. The processor microcode that interprets this instruction sends a list hardware command to
15 the data mover instructing it to move the data areas to or from the operating system zone, STEP 520. The data mover receives this second list command, determines that the coupling facility zone is prepared, and moves the data.

When all of the data areas have been transferred, the
20 data mover updates the CBOB to signal the successful completion of data transfer to the coupling facility zone, STEP 522. Thereafter, or if there is no data to be transferred, INQUIRY 516 (FIG. 5a), the coupling facility zone executes another instruction to move the message
25 response block (MRB), STEP 524. The microcode that interprets this instruction sends a respond hardware command to the data mover instructing it to move the MRB from CBOB 422 to MCB/MRB area 402 in the operating system zone, STEP

526. One example of a respond command is described in Elko et al., U.S. Patent No. 5,706,432, entitled "Mechanism For Receiving Messages At A Coupling Facility", issued January 6, 1998, which is hereby incorporated herein by reference in its entirety.

When the MRB is received, the data mover either updates LCB 404 or interrupts a processor to indicate the successful completion of the message, STEP 528. This completes the data flow of a primary message in which data is moved from an operating system zone to a coupling facility zone. Additional details relating to the processing of primary messages is described further below with reference to FIGs. 10a-10b.

In the logic described above, various steps may be performed in differing order. For instance, the order in which the list commands are received by the data mover may be in a different order than described above. Other steps may also be varied without departing from the spirit of the present invention. The above is only one embodiment.

One embodiment of the flow of data when a message (i.e., a secondary message) is passed from a coupling facility zone to an operating system zone is described with reference to FIGs. 4 and 6. In one example, when a secondary message is to be sent from coupling facility zone 420 (FIG. 4) to operating system zone 400, the coupling facility zone puts the command (MCB) into CBOB 434 and executes an instruction to move the command, STEP 600 (FIG. 6). The processor

microcode that interprets the command issues a send
secondary hardware command to the data mover instructing it
to move the MCB from the CBOB directly to LCB 412, STEP 602.
Then, a processor (any of the processors in the CPC) is
5 interrupted to process the command, STEP 604.

When the command is finished, and the response (MRB) is
placed in LCB 412, STEP 606, a respond hardware command is
sent to the data mover instructing it to send the MRB from
the LCB to CBOB 434 in the coupling facility zone's main
10 memory, STEP 608. After the data mover successfully moves
the MRB, it updates the CBOB to signal completion of the
command to the coupling facility zone, STEP 610. This
completes the data flow of a secondary message.

Further details regarding the movement of data and an
15 implementation of the data mover are described with
reference to FIGs. 7-11. Referring to FIG. 7, data mover
208 includes, for instance, a queued memory line mover 700
and control logic 702, each of which is described in further
detail below.

20 In one embodiment, queued memory line mover 700 receives
from control logic 702 queue entries 704, each of which
represents, for instance, two memory requests (e.g., a fetch
and a store of a command (MCB), a response (MRB), or data);
and immediate requests 706, which represents high priority
25 requests (e.g., requests for addresses) that are not queued.
Additionally, queued memory line mover 700 sends status

information 707 to the control logic indicating completion of the memory requests.

Queued memory line mover 700 is the interface to main memory 708, and in accordance with an aspect of the present invention, is responsible for moving the data between zones in main memory. In one example, the queued memory line mover moves data from one zone's main memory to another zone's main memory in groups of memory lines. Each memory line is typically 128 bytes and is on address boundaries of 128 bytes, but larger and smaller lines can be used. The memory requests to main memory subsystem 708 instruct it to either fetch some number of data bytes (up to, e.g., 128 bytes, a memory line) or store some number of bytes (also up to, e.g., 128 bytes). Further details of queued memory line mover 700 are described with reference to FIG. 8.

Referring to FIG. 8, queued memory line mover 700 includes, for instance, a FIFO (first in, first out) array 800 that temporarily holds information received from control logic 702 to form memory requests; a fetch state machine 802 that forms memory fetch requests to fetch data from a main memory of one zone; a store state machine 804 that forms memory store requests to store the fetched data in a main memory of another zone; and one or more line buffers 806 that temporarily store the fetched data. Each of the components of queued memory line mover 700 is further described below.

FIFO array 800 includes a plurality of queue entries 808, and each queue entry 808 includes, for example, a main memory source address 810 indicating the location in main memory of the requested data; a destination address 812 5 indicating the location in main memory to store the retrieved data; a length parameter (e.g., byte count) 814 indicating the size of the data; and an ID 815 identifying the queue entry. Source memory address 810 and destination memory address 812 each have two parts. The first is the 10 zone ID and the second is the address within the zone; each zone's addressing starts at address zero, in this embodiment.

Fetch state machine 802 reads source addresses 810 and length parameters 814 from FIFO queue 800, and forms memory 15 fetch requests 816 to the main memory subsystem. The memory fetch request includes, for instance, an indication that it is a memory request, an address from where it wants the data (i.e., source address 810), and a count of the data (i.e., length 814). When the data is returned (818) from the main 20 memory subsystem, it is temporarily stored in one of line buffers 806.

The line buffers hold the fetched data until store operations are processed. In addition to the fetched data, each line buffer also includes control information 826 used 25 in associating the data with a queue entry, as described below.

Main memory subsystem 708 is capable of processing several memory requests at once. In other words, multiple memory requests can be outstanding simultaneously, and the responses to those requests may not necessarily be in the
5 same order as the initial requests. This improves the throughput of the main memory subsystem. For the queued memory line mover to exploit this feature of the main memory subsystem, multiple line buffers 806 are used. (The number of line buffers depends on the number of outstanding
10 requests allowed.) Each time a memory fetch operation is sent to the main memory subsystem, a line buffer is allocated. If all of the line buffers are busy (waiting for their associated memory store operations to finish), no new memory fetch requests are made. Also, when any line buffer
15 becomes full, the queued memory line mover initiates a memory store operation.

When a line buffer is full, store state machine 804 reads destination addresses 812 and length parameters 814 from FIFO queue 800, and forms memory store requests 820 to
20 the main memory subsystem. In this example, the data to be stored is included with the request and thus, would be included with command 820. However, in another example, the data may be separate from the command, and thus, forwarded on data line 822. When the memory store operation
25 completes, the main memory subsystem returns status to the queued memory line mover, which forwards the status onto control logic 702.

One embodiment of control logic 702 is described with reference to FIGs. 7 and 9a-9b. Referring initially to FIG. 7, control logic 702 receives commands (e.g., hardware commands) from one or more processors 710 of one or more of 5 the zones, and sends interrupts to the processors. Further, it generates queue entries 704, which are provided to queued memory line mover 700; it updates control blocks (such as LCBs and CBOBs); and it fetches addresses for data movement over immediate requests path 706.

10 In one example, an immediate request indicates a high priority request, and thus, does not use the array or line buffers of the queued memory line mover. Instead, the addresses, data length, and data of the immediate request are received by the queued memory line mover directly from 15 the control logic, and that memory request is executed immediately by the queued memory line mover. In particular, when the control logic wishes to fetch address information from main memory, it sends the address of the address field in a control block along with a byte count (typically eight 20 bytes) to the queued memory line mover. The queued memory line mover then sends a fetch command to the main memory subsystem, when it is allowed to send the next memory request. When the data is returned, it is sent back to the control logic along with a status code.

25 When the control logic wishes to update an LCB or CBOB, it sends the address of the field within the LCB or CBOB and the data (typically eight bytes) to the queued memory line mover. The queued memory line mover then sends a store

request to the main memory subsystem, when it is allowed to send the next memory request. When the memory operation completes, the queued memory line mover sends a completion status back to the control logic.

5 Further details of control logic 702 are described with reference to FIG. 9a. Control logic 702 receives commands (e.g., hardware commands, such as list, send, etc., as described herein) from one or more processors 710 and sends interrupts back to the processors. Additionally, the
10 control logic sends queue entries 704 and immediate memory requests 706 to queued memory line mover 700. In one embodiment, the control logic is divided into two parts: one part 900a for processing primary messages from an operating system zone to a coupling facility zone, and the other part
15 900b for processing secondary messages from a coupling facility zone to an operating system zone.

The control logic for both the primary and secondary messages includes, for instance, an address register for each of the memory areas shown in FIG. 4. For example,
20 MCB/MRB 402 is stored in MCB/MRB address register 906 in the control logic; LCB 404 corresponds to LCB 908; address list 410 corresponds to address list 910; data 406, 408 correspond to data 912; CBOB 422 corresponds to CBOB 914; address list 432 corresponds to address list 916; data 424,
25 426, 428, 430 corresponds to data 918; LCB 412 corresponds to LCB 920; and CBOB 434 corresponds to CBOB 922.

Primary message address registers 906, 908, 910, 912 belong to operating system zones and primary address registers 914, 916, 918 belong to coupling facility zones. Likewise, secondary message address register 920 belongs to the operating system zones, and secondary address register 922 belongs to the coupling facility zones. There are also parameter registers 924, 926 and status registers 928, 930. Parameter registers contain message control parameters such as data transfer direction indication, MCB length, number of data areas, data area length(s), MRB length, etc. The status registers indicate the progress of the message. For example, state information in the status register indicates the prepared state of the message, an offset into the address list, the current (or next) memory line transfer sent to the queued memory line mover, etc.

In the example described herein, each of the registers is shown as a group of registers (e.g., three registers) indicating the capability of processing multiple messages simultaneously. The facilities to exchange a single message are collectively called a message buffer. In practice, the groups may contain any number of registers depending on the level of message multiplexing desired by the CPC.

When the control logic wishes to move a part of a message (i.e., a message segment, such as an MCB, an MRB, or data) between zones, it generates one or more queue entries to send to the queued memory line mover. For example, consider transferring a 264 byte MCB. Assuming a memory line size of 128 bytes and assuming that the MCB starts on a

POU919980103US1

256 byte main memory address boundary, the control logic generates three queue entries. The first entry has the source and destination addresses and a length of 128 bytes, the second entry has the source (+128) and destination 5 (+128) addresses and a length of 128 bytes, and the third entry has the source (+256) and destination (+256) addresses and a length of 8 bytes. Each entry also includes an identifier for that entry.

The control logic keeps track of the progress of each 10 queue entry. Since the queue entries may complete in a different order than they were queued, the control logic makes sure that all of the queue entries for a particular part of the message are completed before proceeding. The control logic also interleaves queue entries for multiple 15 simultaneous messages. This technique gives fairness while ensuring forward progress, and is further described with reference to FIGs. 7-9b.

In particular, each queue entry sent from the control logic to the FIFO of the queued memory line mover includes 20 ID 815 (FIG. 8). When there is a pending request in the FIFO (each entry represents one or more requests (i.e., a fetch and/or a store)) and a line buffer is available, the queued memory line mover takes the request from the FIFO and stores the FIFO location in controls 826 of the line buffer. 25 The queued memory line mover then makes a fetch request 816 to the main memory subsystem. When the fetch request completes and the data line is in the line buffer, the FIFO location information is used to index back into the FIFO to

POU919980103US1

find the destination address for the memory line of data. Further, (e.g., at substantially the same time), the FIFO location in line buffer controls 826 is replaced by the ID from the FIFO, and the FIFO location is now free to receive
5 the next request from the control logic. Next, the queued memory line mover makes a store request 820 to the main memory subsystem. When the store request completes successfully, the queued memory line mover forwards the memory response with the ID field in line buffer controls
10 826 back to the control logic.

Since it typically requires several line move operations to move each part of a message (i.e., a message segment), the control logic keeps a table 950 (FIG. 9b) that is associated with the active operations. For example, one
15 implementation allows four message segments (i.e., different segments, the same segments, or a combination thereof) to each have up to four requests outstanding to the queued memory line mover. That is, each row 952 of the table represents a message segment, and each column of the first
20 four columns represents a memory request of that segment. The last column identifies the message segment.

As message segments are processed, they are assigned to a group of four bits to keep track of the progress of the segment, and a message segment identifier is placed in table
25 950. Each subsequent request sent to the queued memory line mover is given a four bit ID 815. Two of the bits identify which of the groups of four bits (i.e., a row) is being used to process the message segment, and the other two bits

identify one of the four bits in the group (i.e., a column). When four requests are made, all four bits of the group are set. As the queued memory line mover finishes each request, the ID field in the response sent back to the control logic
5 resets the appropriate bit in table 950, and the control logic is free to make the next request.

The priority scheme used by the control logic is first come first served at the message segment level. Once a segment is chosen, one of the four bit groups of table 950
10 is assigned to the message segment. All of the memory lines are moved for the message segment before the group is reassigned to the next message segment.

The table described above is only one example. In other embodiments, the table may have more or less rows and/or
15 columns, and may contain more, less or different information.

Further details regarding an embodiment of the operation of control logic 702 is described with reference to FIGS. 10-11. In particular, the processing of a primary message
20 is further described with reference to FIGs. 10a and 10b, and the processing of a secondary message is further described with reference to FIG. 11.

Referring to FIG. 10a, when an operating system zone wishes to send a primary message to a coupling facility
25 zone, it builds an MCB in the MCB/MRB area and executes a send message instruction, STEP 1000. The microcode

executing the send message instruction copies the MCB to the LCB where the operating system zone ID is written into the MCB header, STEP 1002. The microcode then sends a start hardware command to the data mover, STEP 1004. This
5 hardware command includes the address of the MCB/MRB area in the operating system zone's main memory.

When the control logic receives the start hardware command, it checks the prepared state of the coupling facility zone, INQUIRY 1006. In particular, the control
10 logic checks a prepared field in a status buffer 928 to determine whether the coupling facility zone is prepared. If the coupling facility zone is not prepared, the control logic generates an interrupt to the processor, STEP 1008, and does not transfer the MCB to the coupling facility zone.
15 The interrupt causes the processor to interrogate the control logic to determine that an error response should be sent back to the operating system zone, STEP 1010. The error response hardware command is sent to the data mover instructing it to send the error response MRB from a
20 location in HSA to the MCB/MRB area in the operating system zone's main memory, STEP 1012. Note that the MCB/MRB area address is already known to the data mover by the previous start hardware command. The control logic sends the response to the operating system zone and signals/interrupts
25 the operating system zone, STEP 1014. This completes processing of the primary message, STEP 1018.

Returning to INQUIRY 1006, if the coupling facility zone is prepared, the control logic generates queue entries to

POU919980103US1

move the MCB from the LCB in the operating system zone to the CBOB in the coupling facility zone, STEP 1020. The queue entries are processed by the queued memory line mover. When the control logic detects that the entire MCB has been
5 moved to the CBOB, it sets a signal (in this embodiment through a busy bit in the hardware observable to the coupling facility zone) informing the coupling facility zone that a message has arrived.

Thereafter, a determination is made as to whether there
10 are data areas to be moved, INQUIRY 1022. If there are data areas to be moved, the logic of FIG. 10b is performed, as described below. When there is no data to transfer, INQUIRY 1022, or after the data has been successfully transferred, STEP 1024, the coupling facility detects this condition and
15 generates a response, STEP 1026. In particular, the coupling facility microcode sends a respond hardware command to the data mover instructing it to generate memory requests to move the MRB from the CBOB in the coupling facility zone's main memory to the MCB/MRB area in the operating
20 system zone's memory. The control logic generates the memory requests to move the MRB to the MCB/MRB area, and signals/interrupts the operating system zone, STEP 1028. This completes processing of the primary message, STEP 1030.

Returning to INQUIRY 1022, if there is data to be
25 transferred, the logic of FIG. 10b is employed. In one example, initially, the operating system zone prepares a list of data addresses, and then the operating system microcode sends a list hardware command to the data mover,

STEP 1050. This command includes the direction of data transfer (e.g., read/write), the address of the beginning of the address list, the length of each data area, and the number of data areas to transfer. The data mover does not
5 start to transfer the data until microcode in the coupling facility zone also sends a list hardware command to the data mover, STEP 1052. This list hardware command from the coupling facility zone also includes the direction of data transfer, the address of the beginning of the address list,
10 the length of each data area, and the number of data areas to transfer. At this point, the control logic has the information required to fetch the first address in each of the two lists, STEP 1054.

After the addresses have been fetched, the control logic
15 determines if the data transfer is for a read or write operation, INQUIRY 1056. If it is a read operation, the data transfer is from the coupling facility zone to the operating system zone. In this case, the addresses from the coupling facility zone are placed in the source address area
20 of the queue entries and the addresses from the operating system zone are placed in the destination address area of the queue entries, STEP 1058. However, if it is a write operation, the data transfer is from the operating system zone to the coupling facility zone. In this case, the
25 addresses from the operating system zone are placed in the source address area of the queue entries and the addresses from the coupling facility zone are placed in the destination address area of the queue entries, STEP 1060.

After all of the data for the first address of the list has been transferred, a determination is made as to whether the data transfer is complete, INQUIRY 1062. If it is not complete, the control logic fetches the next address from
5 the list, STEP 1054. Note that the lengths of each data area (pointed to by an address in the list) do not have to be the same in the operating system and coupling facility zones. The control logic fetches addresses from the operating system and coupling facility lists independently
10 after their respective byte counts have been exhausted.

Returning to INQUIRY 1062, when the data transfer is complete, the control logic informs the coupling facility zone by updating the CBOB, STEP 1064. Control is then passed back to FIG. 10a, where the message response is sent,
15 STEP 1066.

One embodiment of the flow for processing secondary messages is described with reference to FIG. 11. Initially, the coupling facility issues a send secondary message, STEP 1100. Then, the microcode issues a send hardware command to
20 the data mover that includes the length of the secondary message, STEP 1102. Therafter, a determination is made as to whether the other end is prepared, INQUIRY 1104. If it is not prepared, an error is signaled, STEP 1106. Otherwise, the control logic generates memory requests to
25 send the MCB from the CBOB in the coupling facility zone's main memory to the LCB in the operating system zone's memory, STEP 1108. The operating system zone's microcode detects the secondary message, processes the command, and

generates a response, STEP 1110. The operating system microcode then sends a respond hardware command to the data mover that includes the length of the response, STEP 1112. The control logic moves the response MRB from the LCB in the
5 operating system zone's main memory to the CBOB in the coupling facility zone's memory completing the message exchange, STEP 1114.

Described in detail above is a data mover used to move data (e.g., commands, responses and/or data areas) between
10 zones of a central processing complex. Advantageously, the data mover moves this data without using processor instructions and without using a channel interface. This reduces the cost of the data movement, and optimizes the movement, since the hardware employed is specially tailored
15 to directly move the data between the zones. The data mover also provides memory protection, in a similar manner as channels.

The flow diagrams depicted herein are just examples. There may be many variations to these diagrams or the steps
20 (or operations) described therein without departing from the spirit of the invention. For instance, the steps may be performed in a differing order, or steps may be added, deleted or modified. All of these variations are considered a part of the claimed invention.

25 Although preferred embodiments have been depicted and described in detail herein, it will be apparent to those skilled in the relevant art that various modifications,

additions, substitutions and the like can be made without departing from the spirit of the invention and these are therefore considered to be within the scope of the invention as defined in the following claims.

0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232